

# Feature Selection for Regression Problems Based on the Morisita Estimator of Intrinsic Dimension

Jean GOLAY, Michael LEUENBERGER and Mikhail KANEVSKI

*Institute of Earth Surface Dynamics, Faculty of Geosciences and Environment,  
University of Lausanne, 1015 Lausanne, Switzerland. Email: jean.golay@unil.ch.*

---

## Abstract

Data acquisition, storage and management have been improved, while the key factors of many phenomena are not well known. Consequently, irrelevant and redundant features artificially increase the size of datasets, which complicates learning tasks, such as regression. To address this problem, feature selection methods have been proposed. This paper introduces a new supervised filter based on the Morisita estimator of intrinsic dimension. It can identify relevant features and distinguish between redundant and irrelevant information. Besides, it offers a clear graphical representation of the results and it can be easily implemented in different programming languages. Comprehensive numerical experiments are conducted using simulated datasets characterized by different levels of complexity, sample size and noise. The suggested algorithm is also successfully tested on a selection of real world applications and compared with RReliefF using extreme learning machine. In addition, a new measure of feature relevance is presented and discussed.

*Keywords:* Feature selection, Intrinsic dimension, Morisita index, Measure of relevance, Data mining

---

## 1. Introduction

In data mining, it is often not known a priori what features (or input variables <sup>1</sup>) are truly necessary to capture the main characteristics of a studied phenomenon. This lack of knowledge implies that many of the considered features are irrelevant or redundant. They artificially increase the

---

<sup>1</sup>In this paper, the term “feature” is used as a synonym for “input variable”.  
*Preprint submitted to Elsevier*

dimension  $E$  of the Euclidean space in which the data points are embedded ( $E$  equals the number of input and output variables under consideration). This is a serious matter, since fast improvements in data acquisition, storage and management cause the number of redundant and irrelevant features to increase. As a consequence, the interpretation of the results becomes more complicated and, unless the sample size  $N$  grows exponentially with  $E$ , the curse of dimensionality [1] may reduce the overall accuracy yielded by any learning algorithm. Besides, large  $N$  and  $E$  are also difficult to deal with because of computer performance limitations.

In regression and classification, these issues are often addressed by implementing supervised feature selection methods [2–5]. They can be broadly subdivided into filter (e.g. RReliefF [6], mRMR [7] and CFS [8]), wrapper [9, 10] and embedded methods (e.g. the Lasso [11] and random forest [12]). Filters rank features or subsets of features according to a relevance measure independently of any predictive model, while wrappers use an evaluation criterion involving a learning machine. Both approaches can be used with search strategies, since an exhaustive exploration of the  $2^{\#Feat.} - 1$  models (all the possible combinations of features) is often computationally infeasible. Greedy strategies [13, 14], such as Sequential Forward Selection (SFS), can be distinguished from stochastic ones (e.g. simulated annealing [15, 16] and particle swarm optimization [17, 18]). Regarding the embedded methods, the feature selection is a by-product of a training procedure. It can be achieved by the addition of constraints in the cost function of a predictive model (e.g. the Lasso [11] and ridge regression [19]) or it can be more specific to a given algorithm (e.g. random forest [12] and adaptive general regression neural networks [20, 21]).

The present paper<sup>2</sup> deals with a new SFS filter algorithm. It relies on the idea that, although data points are embedded in  $E$ -dimensional spaces, they often reside on lower  $M$ -dimensional manifolds [23–25]. The value  $M$  ( $\leq E$ ) is called Intrinsic Dimension (ID) and it can be estimated using the Morisita estimator of ID, introduced in [26], which is closely related to the fractal theory. The proposed filter algorithm is supervised, designed for regression problems and based on this new ID estimator. It also keeps the simplicity of the Fractal Dimension Reduction (FDR) algorithm introduced

---

<sup>2</sup>The main idea of this paper was partly presented at the 23rd symposium on artificial neural networks, computational intelligence and machine learning (ESANN2015) [22].




Figure 1: 50 values were sampled from two uniformly distributed variables  $V_1$  and  $V_2$ . The two variables are (a) independent, (b) linearly dependent or (c) non-linearly dependent.

in [27]. Finally, the results show the ability of this new filter to capture non-linear relationships and to effectively identify both redundant and irrelevant information.

The paper is organized as follows. Section 2 reviews previous work on ID-based feature selection approaches. The Morisita estimator of ID is shortly presented in Section 3 (for the completeness of the paper). Section 4 introduces the Morisita-based filter and Section 5 is devoted to numerical experiments conducted on simulated data of varying complexity. In Section 6, real world applications from publicly accessible repositories are presented and a comparison with a benchmark algorithm, RReliefF [6], is carried out using Extreme Learning Machine (ELM) [28]. Finally, conclusions are drawn in the last section with a special emphasis on future challenges and applications.

## 2. Related Work

The concept of ID can be extended to the more general case where the data ID may be a non-integer dimension  $D$  [24, 27, 29]. The value  $D$  is estimated by using fractal-based methods which have been presented in [24, 25, 30] and successfully implemented in various fields, such as physics [31], cosmology [32], meteorology [33] and pattern recognition [34]. These methods rely on well-known fractal dimensions (e.g. the box-counting dimension [35, 36], the correlation dimension [31] and Rényi's dimensions of  $q$ th order [37]) and they can be used in feature selection [27, 38] and dimensionality reduction [24] to detect dependencies between variables (or features).

Traina et al. [27, 39] have opened up new prospects for the effective use of ID estimation in data mining by introducing the Fractal Dimension Reduction (FDR) algorithm. FDR executes an unsupervised procedure of feature selection aiming to remove from a dataset all the redundant variables. The fundamental idea is that fully redundant variables do not contribute to the value of the data ID.

This idea is illustrated in Figure 1 where 50 points were sampled from two uniformly distributed variables  $V_1$  and  $V_2$ . In the first panel,  $V_1$  and  $V_2$  are independent, which means that they are not redundant, and one has that:

$$ID(V_1, V_2) \approx ID(V_1) + ID(V_2) \approx 1 + 1 = 2 \quad (1)$$

where  $ID(\cdot)$  denotes the ID of a dataset. It indicates that both  $V_1$  and  $V_2$  contribute to increasing the value of  $ID(V_1, V_2)$  by about 1, which is, by construction, equal to the ID of each variable (i.e.  $ID(V_1)$  and  $ID(V_2)$ ). Conversely, the removal of either  $V_1$  or  $V_2$  would lead to a reduction in the data ID from about 2 (i.e. the dimension of the embedding space) to 1 (i.e. the ID of a single variable) and information would be irreparably lost. In contrast, the other two panels show what happens when  $V_1$  and  $V_2$  are fully redundant with each other. One has that:

$$ID(V_1, V_2) \approx ID(V_1) \approx ID(V_2) \approx 1 \quad (2)$$

where the ID of the full dataset is approximately equal to the topological dimension of a smooth line. This means that the contribution of only one variable is enough to reach the value of  $ID(V_1, V_2)$  and the remaining one can be disregarded without losing any information.

Based on these considerations, the FDR algorithm removes the redundant variables from a dataset by implementing a Sequential Backward Selection (SBS) strategy [13]. Besides, it uses Rényi's dimension of order  $q = 2$ ,  $D_2$ , for the ID estimation. Following the same principles, De Sousa et al. [40] examined additional developments to FDR and presented a new algorithm for identifying subgroups of correlated variables.

FDR is designed to carry out unsupervised tasks and it is not able to distinguish between variables that are relevant to a learning process and those that are irrelevant. The reason is that such variables can all contribute to

the data ID. For instance, in the left-hand panel of Figure 1,  $V_1$  could be regarded as irrelevant to the learning of  $V_2$ , but it would be selected by FDR because it makes the data ID increase by about 1. Consequently, different studies were carried out to adapt FDR to supervised learning. Lee et al. [41] suggested decoupling the relevance and redundancy analysis. Following the same idea, Pham et al. [42] used mutual information to identify irrelevant features and combined the results with those of FDR. Finally, Mo and Huang [38] developed an advanced algorithm to detect both redundant and irrelevant information in a single step. This algorithm follows a backward search strategy and relies on the correlation dimension,  $df_{cor}$ , for the estimation of the data ID.

The filter algorithm suggested in this paper is designed, in such a way that it combines the advantages of both FDR and Mo’s algorithm: it can deal with non-linear dependencies, it does not rely on any user-defined threshold, it can discriminate between redundant and irrelevant information and the results can be easily summarized in informative plots. Moreover, it can cope with high-dimensional datasets thanks to its SFS search strategy and it uses the Morisita estimator of ID which was shown to yield comparable or better results than  $D_2$  and  $df_{cor}$  [26].

### 3. The Morisita Estimator of Intrinsic Dimension

The Morisita estimator of ID,  $M_m$ , has been recently introduced [26]. It is a fractal-based ID estimator derived from the multipoint Morisita index  $I_{m,\delta}$  [30, 44] (named after Masaaki Morisita who proposed the first version of the index to study the spatial clustering of ecological data [45]).  $I_{m,\delta}$  is computed by superimposing a grid of  $Q$  quadrats of diagonal size  $\delta$  onto the data points (see Figure 2). It measures how many times more likely it is that  $m$  ( $m \geq 2$ ) points selected at random will be from the same quadrat than it would be if all the  $N$  points of the studied dataset were distributed according to a random distribution generated from a Poisson process (i.e. complete spatial randomness). The formula is the following:

$$I_{m,\delta} = Q^{m-1} \frac{\sum_{i=1}^Q n_i(n_i - 1)(n_i - 2) \cdots (n_i - m + 1)}{N(N - 1)(N - 2) \cdots (N - m + 1)} \quad (3)$$

where  $n_i$  is the number of points in the  $i^{th}$  quadrat. For a fixed value of  $m$ ,  $I_{m,\delta}$  is calculated for a chosen scale range. If a dataset approximates a

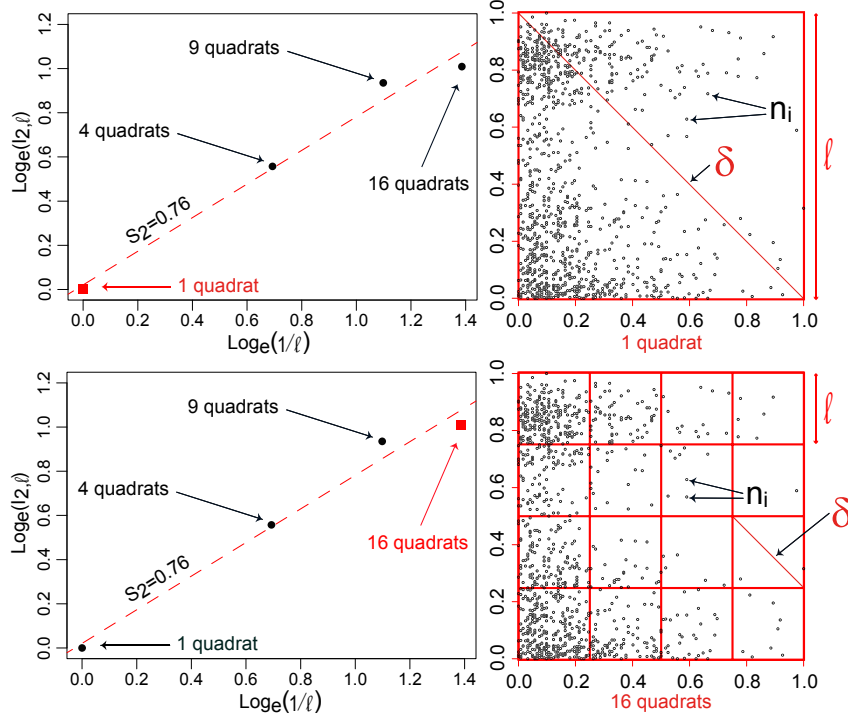


Figure 2: Illustration of the way the Morisita slope,  $S_2$ , is computed for a 2-dimensional problem. In the two left-hand panels,  $S_2$  is the slope of the dashed line and the red squares correspond to the values of  $\log_e(I_{2,\ell})$  calculated with the grids displayed on the right. The R dataset “Quakes” [43] was used and the data were rescaled to the  $[0, 1]$  interval.

fractal behavior (i.e. is self-similar) within this range, the relationship of the plot relating  $\log(I_{m,\delta})$  to  $\log(1/\delta)$  is linear and the slope of the regression line is defined as the Morisita slope  $S_m$ . Finally,  $M_m$  is expressed as:

$$M_m = E - \left( \frac{S_m}{m-1} \right). \quad (4)$$

In practice, each variable is rescaled to the  $[0, 1]$  interval and  $\delta$  can be replaced with the quadrat edge length,  $\ell$ , with  $\ell^{-1}$  being the number of grid subdivisions (see Figure 2). Then, a set of  $R$  values of  $\ell$  (or  $\ell^{-1}$ ) is chosen so that it captures the linear part of the log-log plot. In the rest of this paper, only  $M_2$  will be used and it will be computed with an algorithm called Morisita INDEX for Intrinsic Dimension estimation (MINDID) [26] whose complexity is  $\mathcal{O}(N * E * R)$ .

#### 4. The Morisita-based Filter for Regression Problems

The Morisita-Based Filter for Regression (MBFR) relies on three observations following from the work by Traina et al. [27], De Sousa et al. [40] and Mo and Huang [38]:

1. Given an output variable  $Y$  generated from  $k$  relevant and non-redundant input variables  $X_1, \dots, X_k$ , one has that:

$$ID(X_1, \dots, X_k, Y) - ID(X_1, \dots, X_k) \approx 0 \quad (5)$$

where  $ID(\cdot)$  denotes the (possibly non-integer) ID of a dataset.

2. Given  $i$  irrelevant input variables  $I_1, \dots, I_i$  completely independent of  $Y$ , one has that:

$$ID(I_1, \dots, I_i, Y) - ID(I_1, \dots, I_i) \approx ID(Y) \quad (6)$$

3. Given a randomly selected subset of  $X_1, \dots, X_k$  of size  $r$  with  $0 < r < k$  and  $k > 1$ ,  $j$  redundant input variables  $J_1, \dots, J_j$  related to some or all of  $X_1, \dots, X_r$  and all the  $i$  irrelevant input variables  $I_1, \dots, I_i$ , one has that:

$$ID(X_1, \dots, X_r, J_1, \dots, J_j, I_1, \dots, I_i, Y) - ID(X_1, \dots, X_r, J_1, \dots, J_j, I_1, \dots, I_i) \approx H \quad (7)$$

where  $H \in ]0, ID(Y)[$  and  $H$  decreases to 0 as  $r$  increases to  $k$ .

The difference

$$Diss(F, Y) = ID(F, Y) - ID(F) \quad (8)$$

can thus be suggested as a way of measuring the dissimilarity (i.e the independence) between  $Y$  and a set  $F$  of features (e.g.  $F = \{X_1, \dots, X_k\}$ ), among which only the relevant ones (i.e. the non-redundant features on which  $Y$  depends) contribute to reducing the dissimilarity. Based on that idea, MBFR (see Algorithm 1) aims at retrieving the relevant features available in a dataset by sorting each subset of variables according to its dissimilarity with  $Y$ . MBFR implements a SFS search strategy and relies on the Morisita estimator of ID and the MINDID algorithm [26] to estimate  $Diss$ :

$$\widehat{Diss}(F, Y) = M_2(F, Y) - M_2(F). \quad (9)$$

---

**Algorithm 1** MBFR

---

**INPUT:**

A dataset  $A$  with  $E - 1$  features  $F_{1,\dots,E-1}$  and one output variable  $Y$ .

A vector  $L$  of values  $\ell$ .

An integer  $C$  ( $\leq E - 1$ ) indicating the number of iterations.

Two empty vectors of length  $C$ :  $SelF$  and  $DissF$  for storing, respectively, the name of the selected features and the dissimilarities.

An empty matrix  $Z$  for storing the selected features.

**OUTPUT:**  $SelF$  and  $DissF$ .

- 1: Rescale each feature and  $Y$  to  $[0, 1]$ .
  - 2: **for**  $i = 1$  **to**  $C$  **do**
  - 3:   **for**  $j = 1$  **to**  $(E - i)$  **do**
  - 4:      $\widehat{Diss}(Z, F_j, Y) = M_2(Z, F_j, Y) - M_2(Z, F_j)$  (MINDID used with  $L$ )
  - 5:   **end for**
  - 6:   Store in  $SelF[i]$  the name of the  $F_j$  yielding the lowest value of  $\widehat{Diss}$ .
  - 7:   Store this value of  $\widehat{Diss}$  in  $DissF[i]$ .
  - 8:   Remove the corresponding  $F_j$  from  $A$  and add it into  $Z$ .
  - 9: **end for**
- 

The complexity of the algorithm is  $\mathcal{O}(N * E^2 * R * C)$  where  $C$  ( $\leq E - 1$ ) is the number of iterations of the SFS procedure. For high-dimensional datasets, if  $Diss$  is likely to reach its minimum value after only a few iterations because of many redundant and irrelevant variables,  $C$  can be set to a value lower than  $E - 1$ .

For ease of comparison, the coefficient of dimensional relevance,  $DR(F, Y)$ , can be introduced. It is defined as:

$$DR(F, Y) = 1 - \frac{Diss(F, Y)}{ID(Y)} = 1 - \frac{ID(F, Y) - ID(F)}{ID(Y)} \quad (10)$$

and it can be computed using the Morisita estimator of ID  $M_2$ . In the same way as  $Diss(F, Y)$ ,  $DR(F, Y)$  is able to capture both linear and non-linear relationships between an input and an output space. Besides, it lies between 0 and 1. If the target variable  $Y$  can be completely explained by the considered features  $F$ ,  $DR(F, Y) = 1$ . On the contrary, if all the available features are irrelevant,  $DR(F, Y) = 0$  and, in-between, the closer it is to 1, the greater the explanatory power of  $F$ .



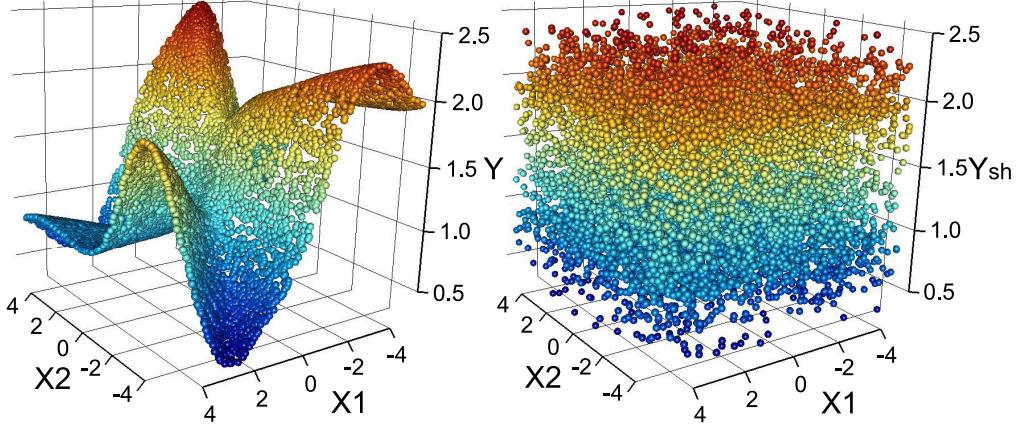


Figure 3: (left) The functional relationship between the output variable  $Y$  and the relevant features  $X_1$  and  $X_2$  of the butterfly dataset; (right) Shuffling of the output variable  $Y$ .

$j$	$\omega_{1,j}$	$\omega_{2,j}$	$\beta_j$
1	0.6655	0.8939	1.3446
2	1.2611	-0.3512	-0.0115
3	0.3961	-1.7827	1.2770
4	-1.7065	-0.5297	0.5962
5	0.8807	1.9574	-0.8530
6	1.8260	0.7962	-0.7290
7	1.3400	1.5001	1.2339
8	1.2919	-0.4462	0.1186
9	-1.3902	1.6856	0.5277
10	0.0743	1.5625	-0.6952

Table 1: Weights used in the construction of the butterfly dataset.

In the rest of this paper, MBFR will be thoroughly tested by using both simulated data and real world applications. It will also be compared with another filter, RReliefF, through the results of different ELM models.

## 5. Experimental Study Using Simulated Data

In this section, the MBFR algorithm is assessed by means of two simulated datasets (see Subsection 5.1) and its overall performance is carefully examined through a set of questions around which the subsections are organized:

- Question 1: How does sample size affect MBFR (see Subsection 5.2)?
- Question 2: How does the complexity of data manifolds affect MBFR (see Subsection 5.3) ?
- Question 3: How can MBFR help distinguish between redundant and irrelevant information (see Subsection 5.4)?
- Question 4: How does MBFR respond to a (partial) lack of relevant information (see Subsection 5.5)?
- Question 5: How does MBFR respond to the presence of noise in data (see Subsection 5.6)?
- Question 6: How better is the coefficient of dimensional relevance compared to a linear measure (see Subsection 5.7)?

Notice also that the R environment [43] was used to implement the MBFR algorithm and to carry out the experiments.

### 5.1. Simulated Datasets

Two simulated datasets were used: the butterfly and the Friedman datasets.

1. The butterfly dataset<sup>3</sup> (see Figure 3): An output variable  $Y$  is generated from two uniformly distributed input variables  $X_1, X_2 \in ]-5, 5[$  by using an Artificial Neural Network (ANN) consisting of one hidden layer of 10 neurons. It can be expressed as:

$$Y = \left[ \sum_{j=1}^{10} \beta_j \text{sig}(X_1 \omega_{1,j} + X_2 \omega_{2,j}) \right] + \varepsilon \quad (11)$$

where  $\omega_{1,j}$  and  $\omega_{2,j}$  are the weights connecting the input variables to the  $j^{th}$  neuron,  $\text{sig}(x) : \mathbb{R} \rightarrow \mathbb{R}$  is a sigmoid transfer function,  $\beta_j$  is the weight between the  $j^{th}$  neuron and the output layer and  $\varepsilon$  is a Gaussian noise with zero mean and varying standard deviation (by default, it is set to 0.00). The exact weights used in the construction of the dataset are given in Table 1. Moreover, the addition of three redundant ( $J$ ) and three irrelevant ( $I$ ) variables is also made to complete the input space:  $J_3 = \log_{10}(X_1 + 5)$ ,  $J_4 = X_1^2 - X_2^2$ ,  $J_5 = X_1^4 - X_2^4$ , a uniformly distributed variable  $I_6 \in ]-5, 5[$ ,  $I_7 = \log_{10}(I_6 + 5)$  and  $I_8 = I_6 + I_7$ . Finally, the butterfly dataset is generated by random sampling of  $X_1$ ,  $X_2$  and  $I_6$ . In this paper, different sample sizes were considered:  $N = 1000, 2000, 10000, 20000$ .

---

<sup>3</sup>It can be downloaded from: <https://sites.google.com/site/jeangolayresearch/>.

2. The Friedman dataset: this dataset uses a function suggested in [46] to test the ability of Multivariate Adaptive Regression Splines (MARS) models to uncover structures in data. The output  $Y$  is given by:

$$Y = 10 \sin(\pi X_1 X_2) + 20(X_3 - 0.5)^2 + 10X_4 + 5X_5 + \varepsilon \quad (12)$$

where  $X_1, X_2, X_3, X_4$  and  $X_5$  are i.i.d. variables following a uniform distribution  $Unif(0, 1)$  and  $\varepsilon$  is a Gaussian random noise with zero mean and unit variance. The input space is then completed by the addition of five irrelevant variables ( $I$ ) following the same uniform distribution:  $I_6, I_7, I_8, I_9$  and  $I_{10} \sim Unif(0, 1)$ . Finally, the Friedman dataset is produced by randomly sampling  $N$  points from the inputs. In this paper, the sample size was set to  $N = 40000$  in accordance with the version of the dataset available on the Regression website [47].

The butterfly and the Friedman datasets are characterized by non-linear structures and their input spaces contain extra variables (i.e. redundant and irrelevant variables) that can be removed without affecting the learning of the target  $Y$ . In the following subsections, MBFR will be subjected to a battery of tests to highlight its ability to select the relevant variables ( $X$ ) and to remove the irrelevant ( $I$ ) and redundant ( $J$ ) ones. Additional experiments will consider shuffled data to examine the response of the algorithm to a complete absence of structure. In parallel, the variability of the results will be examined by means of Monte Carlo simulations: for each experiment, many simulations of the datasets will be generated by repeated random sampling of the input variables.

Notice also that the way the two datasets are constructed leads to the distinction between the data manifolds and the manifolds of the simulated phenomena. The former are built using all the variables (including the output variable), while the latter (referred to as the Friedman and the butterfly manifolds) do not involve the irrelevant features.

Finally, from the perspective of MBFR, a dataset is fully described by the integer values of  $\ell^{-1}$  (see Figure 2). For the butterfly and the Friedman datasets, they were respectively set to  $\{5, 6, \dots, 20\}$  and  $\{1, 2, \dots, 6\}$ . These two sets were chosen, so that, within their bounds, the relationship between  $\log(I_{m,\ell})$  and  $\log(1/\ell)$  was linear. Notice that the upper bound of the

$N$	First Two Features (Occurrences)	$mean(DR)$	$sd(DR)$
1000	$X_1, X_2$ (99); $X_2, X_1$ (1)	0.97	0.02
2000	$X_1, X_2$ (100)	0.97	0.02
10000	$X_1, X_2$ (100)	0.97	0.01

Table 2: The first two features selected by MBFR when applied successively to 100 simulations of the butterfly dataset for different sample sizes. The mean value and the standard deviation of  $DR$  are also provided.

second set is lower than that of the first one. This partially follows from the fact that the Friedman dataset has the greatest ID causing the data points to be sparsely distributed inside the embedding space. As a consequence, beyond  $\ell^{-1} = 5$ , the probability of drawing two points from the same cell is rather low, while it is possible to use values of  $\ell^{-1}$  up to 20 in the case of the butterfly dataset.

### 5.2. Sample Size

MBFR was applied to the butterfly dataset. Three sample sizes were successively considered ( $N = 1000, 2000, 10000$ ) and for each of them, 100 simulations of the data were produced.

Table 2 gives the first two input variables selected by the algorithm and indicates the number of times they were selected first over the simulations. Regardless of the sample size, MBFR always identifies  $X_1$  and  $X_2$  as the most relevant features, although their order can be reversed for  $N = 1000$ . Besides, the explanatory power of these two features was assessed using the coefficient of dimensional relevance  $DR$ . Table 2 provides the mean values of  $DR$  over the simulations as well as the corresponding standard deviations (sd). The means are close to 1, which implies that  $X_1$  and  $X_2$  convey most of the information contained in the dataset and this is in agreement with the data construction:  $X_1$  and  $X_2$  are sufficient to explain  $Y$ , whereas the other features are not necessary or even useless. Moreover, the SFS search strategy enables MBFR to select the most relevant features by exploring a rather low-dimensional space. Consequently, the variability of  $DR$  remains roughly constant over the sample sizes and the standard deviations do not exceed 0.02.

To explore further the potential of the MBFR algorithm, a new series of 100 simulations was generated. For ease of comparison with the next subsections, a constraint was imposed that for each simulation the redundant

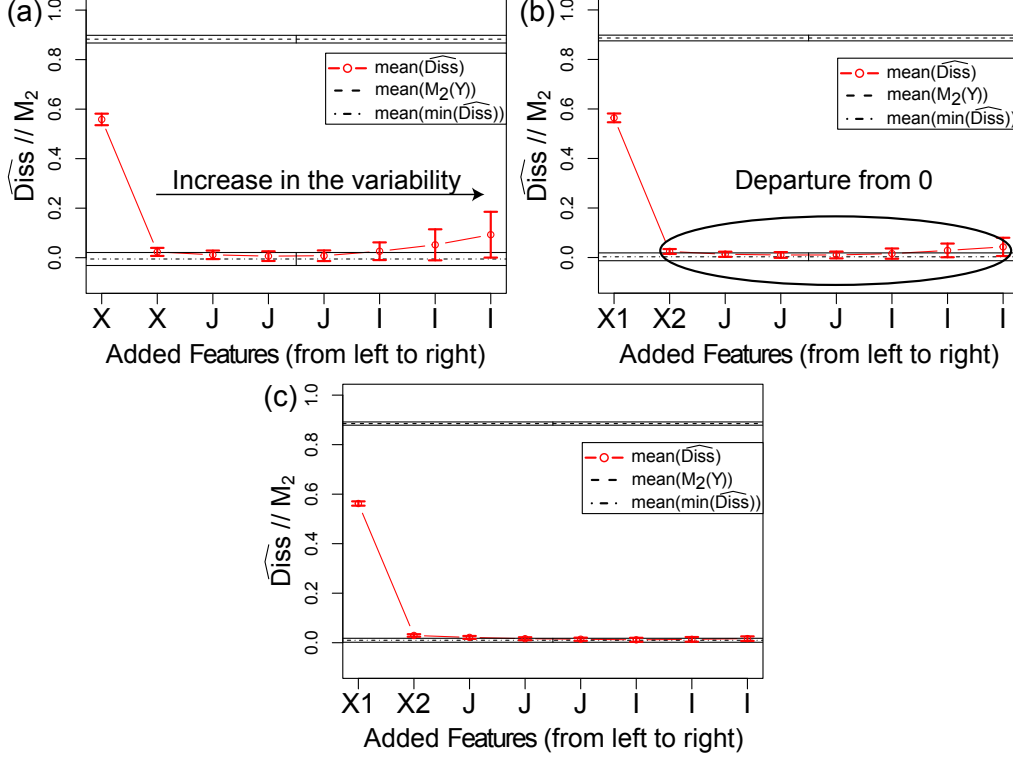


Figure 4: Results of the application of MBFR to 100 simulations of the butterfly data with (a)  $N = 1000$ , (b)  $N = 2000$  and (c)  $N = 10000$ .

variables ( $J$ ) had to be selected by MBFR before the irrelevant ones ( $I$ ). The results are plotted in Figure 4.

The red dots indicate the mean dissimilarity values computed over the simulations by adding to the input space the features appearing on the horizontal axis. In addition, the red bars correspond to the standard deviations. The features are progressively selected from left to right according to the SFS search strategy of MBFR and the names of the redundant and irrelevant ones were shortened to the letters  $J$  and  $I$  because they happened to switch position between the simulations. For the same reasons,  $X_1$  and  $X_2$  were replaced with the letter  $X$  for  $N = 1000$ . Furthermore, in each plot, two additional values are provided: the mean ID estimate of the target variable  $Y$  (i.e.  $mean(M_2(Y))$ ) and the mean of the minimum dissimilarity (i.e.  $mean(\min(\widehat{Diss}))$ ). Their standard deviations are given by the black stripes.

For each sample size,  $X_1$  and  $X_2$  are easily identified as the two relevant features, since they contribute to reducing the dissimilarity from  $M_2(Y)$  to about 0 and clear cut-off points are visible. However, as the number of points is reduced, the variability of the dissimilarity estimates increases. It does not question the potential of the algorithm for feature selection, but it emphasizes two aspects of its implementation: (1) the progressive increase in the variability as more feature are selected and (2) the departure from 0 of the mean dissimilarity estimates after the addition of the second relevant feature (see Figure 4). These two aspects will be addressed in the next subsection.

### 5.3. Complexity of data manifolds

Lower sample sizes highlight that the variability of the dissimilarity values progressively increase as more features are selected (see Figure 4). This response of MBFR is partly due to the presence of the relevant and irrelevant features which amplifies the data ID during the SFS procedure. But it is also related to their non-linear constructions that affect the ID estimates by altering the point clustering on the data manifold.

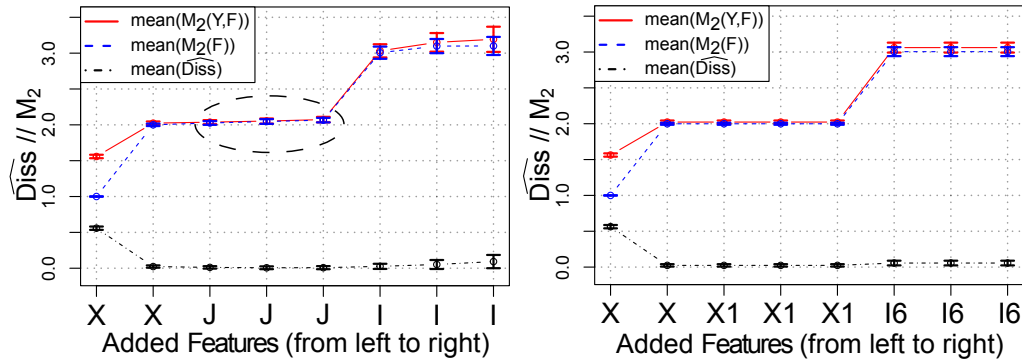


Figure 5: The two panels display the results of each term of Equation 9 for (left) the original butterfly dataset with  $N = 1000$  and for (right) a modified version in which the non-linear dependencies between the features were replaced with pure linear ones.

Figure 5 illustrates the impact of the non-linear dependencies on MBFR. The left-hand panel displays the results of each term of Equation 9 for 100 simulations of the original butterfly dataset. It can be clearly seen that all the redundant features, without exception, influence the ID estimates. In

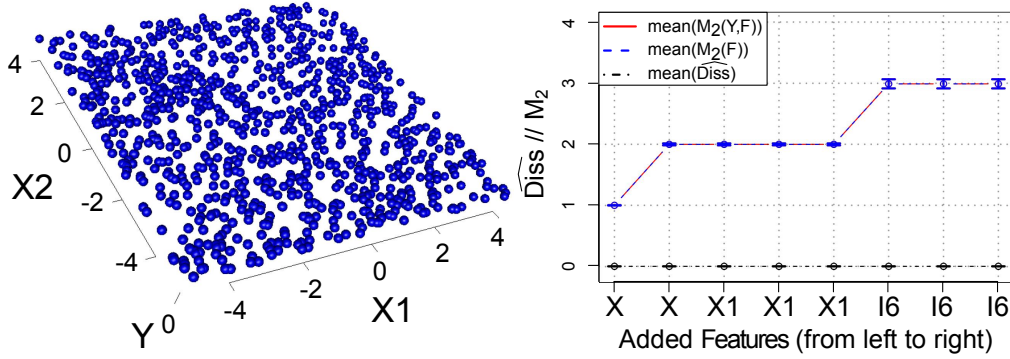


Figure 6: (left) A simplified version of the butterfly manifold for which the variable  $Y$  of the original data was replaced with a zero constant value; (right) the results provided by MBFR when applied to this simplified version of the butterfly dataset for  $N = 1000$ .

contrast, the right-hand panel shows what happens when only pure linear dependencies are considered (i.e.  $J_1$ ,  $J_2$  and  $J_3$  were replaced with  $X_1$ ,  $I_7$  and  $I_8$  were replaced with  $I_6$ ): the mean values and the standard deviations of the ID estimates are modified exclusively by the input variables bringing new information (either useful or useless) and the increase in the variability of the dissimilarity values is no longer progressive. It simply corresponds to  $X_1$ ,  $X_2$  and to the addition the first irrelevant feature.

There is still one aspect of the results of Figure 4 which has not been fully accounted for yet: after the addition of  $X_1$  and  $X_2$ , the dissimilarity values should be equal to 0, but the estimates are slightly higher. Likewise, the mean values of  $DR$  in Table 2 should not be lower than 1. The observed deviations are due to the gap between  $mean(M_2(F, Y))$  and  $mean(M_2(F))$  that is clearly visible in Figure 5 after the addition of the first irrelevant feature. However, if the target variable  $Y$  is replaced with a constant value, the gap vanishes and the dissimilarity estimates equal 0. Figure 6 shows the outcome of MBFR for such a simplified version of the butterfly manifold (again, 100 simulations were used and the non-linear dependencies between the input variables were replaced with pure linear ones for ease of comparison with the right-hand panel of Figure 5). This numerical experiment shows that the shape and the orientation of the data manifold are key factors to explain the gap between the mean ID estimates of Figure 5. The

importance of these factors is related to the quadrats (i.e. the hyper-boxes) of the MINDID algorithm which cannot fit perfectly complex point patterns.

In conclusion, the complexity of the data manifolds (i.e. their shapes, their orientations and the non-linear dependencies between the features) affects the results by altering the terms of Equation 9. However, it does not prevent MBFR from identifying the relevant features.

#### 5.4. *Redundant and Irrelevant Information*

The MBFR algorithm aims to detect the features which are useless (i.e. irrelevant) or not necessary (i.e. redundant) to a regression problem. In addition, it is also able to help distinguish between the two categories of inputs and, more precisely, between redundant and irrelevant information.

For instance, in Figure 5, the first irrelevant feature causes the mean ID estimates to increase by about 1 (i.e. by about  $M_2(I)$ ) and, if it was removed, the second one would have the same effect. In contrast, the redundant features have a much smaller impact which is even hardly noticeable for the relatively low ID values, as highlighted by the dashed ellipse in the left-hand panel. Consequently, the inputs of the butterfly data can be classified as either redundant or irrelevant according to their impacts on the ID estimates.

In real-world applications, a feature (e.g.  $F_1$ ) rejected by MBFR could contain both redundant and irrelevant information. Nevertheless, their exact amounts could still be quantified by using the terms of Equation 9. For instance, if  $F_1$  was partly redundant and partly irrelevant, it would cause an increase in the data ID which would be both higher than 0 (fully redundant) and lower than  $M_2(F_1)$  (fully irrelevant).  $F_1$  would also contain more irrelevant information if the increase was closer to  $M_2(F_1)$  than to 0.

In conclusion, the MBFR algorithm can help distinguish between redundant and irrelevant information by means of the ID estimates on which it relies.

#### 5.5. *Lack of Information*

This subsection investigates the behaviour of the MBFR algorithm when the relevant information is completely or partially missing.






Figure 7: Results of MBFR for (a) the complete Friedman dataset and (b) after the removal of  $X_5$ . In (c) and (d), MBFR was applied to the butterfly dataset after the shuffling of the target variable  $Y$  for, respectively,  $N = 10000$  and  $N = 20000$ .

The top-left panel of Figure 7 displays the results of MBFR achieved for 100 simulations of the Friedman dataset. The algorithm distinctly detects the five relevant features and offers a clear cut-off point. The computations were rerun after the removal of  $X_5$  and the results are given in the top-right panel. This time, the value of  $\min(\widehat{Diss})$  is higher and the difference accounts for the amount of information of  $X_5$ . This last experiment shows that the MBFR algorithm is also able to detect and quantify the absence of relevant features.

It is also worth exploring how MBFR responds to a complete absence of structure between the input and output spaces. A second numerical experiment was set up to that end. The butterfly dataset was used and 100 simulations were generated for  $N = 10000$  and  $N = 20000$ . The target variable  $Y$  of each simulation was then shuffled to destroy the dependencies

between the input and output spaces. The right-hand panel of Figure 3 illustrates the impact of the shuffling procedure on the functional relationship between  $Y$ ,  $X_1$  and  $X_2$ . Finally, MBFR was applied to each simulation and the results are displayed in the bottom panels of Figure 7. As expected, the values of  $\widehat{Diss}$  stay close to  $M_2(Y)$ , which indicates that no features carry relevant information with regard to  $Y$ . The remaining gap between the mean values of  $\widehat{Diss}$  and  $M_2(Y)$  follows from the complexity of the data manifold discussed in Subsection 5.3 and it is more pronounced for the lower sample size, as expected from Subsection 5.2.

Notice also that for comparison purposes, the simulations used in the lower panels of Figure 7 were restricted to those for which the relevant features were selected first, followed successively by the redundant and irrelevant ones. Without this restriction, the final rankings were unpredictable and the irrelevant features could also be picked first. This is consistent with the shuffling procedure which makes all the inputs irrelevant.

In conclusion, the MBFR algorithm provides a way of detecting and quantifying a lack of relevant information. In many cases, a part of that information is not truly missing, but simply corrupted with noise and the next subsection investigates this issue.

#### 5.6. Additional Noise

MBFR should also be able to provide reliable results when the data points are near a manifold instead of being exactly on it. This aspect was investigated by using noisy versions of the butterfly dataset. In the first experiment, the target variable  $Y$  was corrupted with a Gaussian noise (see  $\varepsilon$  in Equation 11) characterized by a zero mean and a standard deviation (sd) of 0.05. 100 simulations of the dataset were generated with  $N = 2000$  and  $N = 10000$ . The results yielded by MBFR are displayed in Figure 8 and a comparison with the default version (i.e. without noise) of the dataset is also presented.

The minimum dissimilarity value is higher once the noise has been added. This was to be expected according to Subsection 5.5, since the noise component partially masks the information that a variable may carry. But nevertheless, the MBFR algorithm is still able to retrieve the relevant variables by providing a clear cut-off point after the addition of the two relevant

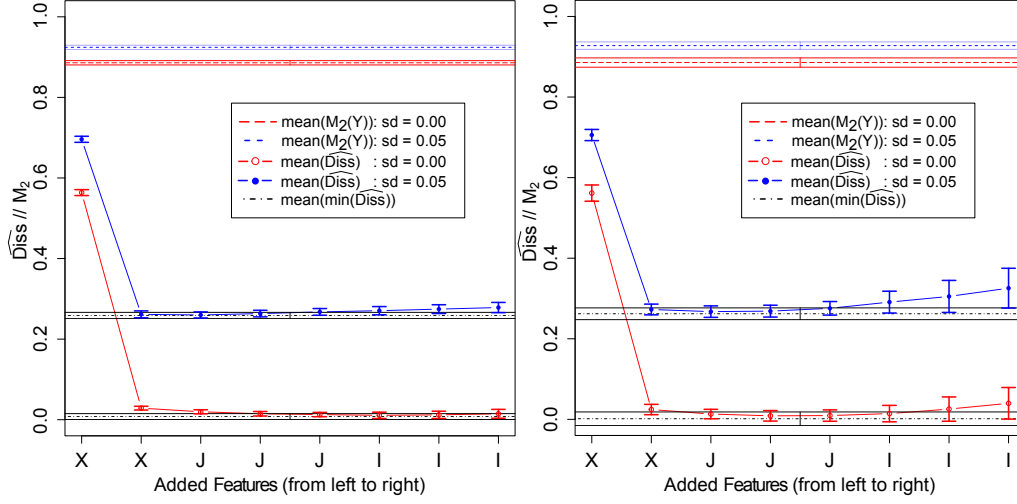


Figure 8: Application of MBFR to two sets of 100 simulations of the butterfly data for (left)  $N = 10000$  and (right)  $N = 2000$ . The first set was not corrupted by noise, while a Gaussian noise was added to the second one.

features.

In Figure 8, the standard deviation of the noise (i.e. 0.05) is about 10% of the original standard deviation of  $Y$  (the mean standard deviation of  $Y$  over 100 simulations is 0.52). To investigate the behaviour of MBFR up to 100% of this value, a second experiment was carried out. The butterfly dataset ( $N = 10000$ ) was generated 100 times for each of the following noise thresholds: 0%, 10%, 20%, 25%, 50%, 75% and 100%. The MBFR algorithm was then applied to each simulation and the first two selected features were recorded at each run. The results are given in Table 3.

The first two features selected by MBFR are given in the middle column, along with the number of times they were selected first over the simulations (in brackets). In addition, the minimum value of  $\widehat{Diss}$ , considering all the simulations, is indicated in the last column. The results show that the first two features can switch position when the noise threshold is higher than 20% and that  $J_3$  is often substituted for  $X_1$ . Nevertheless, even for the greatest noise thresholds, the irrelevant features are never selected first, which means that the MBFR algorithm is still able to recognize the data manifold.

Noise Threshold	First Two Features (Occurrences)	$\min(\widehat{Diss})$
0%	$X_1, X_2$ (100)	0.02
10%	$X_1, X_2$ (100)	0.24
20%	$X_2, J_3$ (2) ; $X_1, X_2$ (4); $X_2, X_1$ (94)	0.46
25%	$X_2, J_3$ (41); $X_2, X_1$ (59)	0.55
50%	$X_2, X_1$ (2) ; $X_2, J_3$ (98)	0.72
75%	$X_2, X_1$ (10); $X_2, J_3$ (90)	0.81
100%	$X_2, X_1$ (10); $X_2, J_3$ (90)	0.85

Table 3: The first two features selected by MBFR when applied successively to 100 simulations of the butterfly dataset for different noise thresholds. The minimum value of  $\widehat{Diss}$  over the simulations is also provided.

### 5.7. The Coefficient of Dimensional Relevance

This subsection aims to compare the coefficient of dimensional relevance  $DR$  (see Equation 10) with the coefficient of determination  $R^2$  derived from linear regression models. To that end, a set of simple scenarios (Sc.) was used, each of which involving three variables from the butterfly dataset:

- Scenario 1:  $X_1$  (input),  $X_2$  (input),  $I_6$  (output).
- Scenario 2:  $X_1$  (input),  $X_2$  (input), the shuffled variable  $Y$  (output).
- Scenario 3:  $X_1$  (input),  $X_2$  (input),  $Y$  (output).
- Scenario 4:  $X_1$  (input),  $X_2$  (input), the variable  $Y$  corrupted with a Gaussian noise  $G \sim \mathcal{N}(0, 0.05^2)$  (output).
- Scenario 5:  $X_1$  (input),  $X_2$  (input), the variable  $Y$  corrupted with a Gaussian noise  $G \sim \mathcal{N}(0, 0.1^2)$  (output).

For each scenario, 100 simulations were produced with  $N = 10000$ . The MBFR algorithm was applied to each of them and  $DR$  was computed. In parallel, linear regression models were fitted to the data using the least-squares approach and the goodness of fit was assessed by  $R^2$ . The results are given in Table 4. When there are no relationships between the inputs and the output (Sc.1 and Sc.2), the two coefficients provide similar results. In contrast, unlike  $R^2$ ,  $DR$  is also able to capture the non-linear shape of the butterfly manifold (Sc.4) and it makes a clear distinction between the default version of the dataset (Sc.4) and the noisy ones (Sc.5 and Sc.6).

## 6. Experimental Study Using Real Data

In this section, the MBFR algorithm is applied to real-world case studies from the UCI machine learning repository [48] and the Regression website

Coefficients	Sc. 1	Sc. 2	Sc. 3	Sc. 4	Sc. 5
$DR$	0.00(0.00)	0.00(0.01)	0.97(0.01)	0.72(0.01)	0.48(0.01)
$R^2$	0.00(0.00)	0.00(0.00)	0.47(0.01)	0.46(0.01)	0.45(0.01)

Table 4: Comparison between  $R^2$  (derived from a linear regression model) and  $DR$ . The value in brackets is the standard deviation over 100 simulations.

[47]. The results are discussed with a special emphasis on the parameter  $\ell^{-1}$ , the coefficient of dimensional relevance  $DR$  and the ability of MBFR to distinguish between redundant and irrelevant information. Finally, a comparison with a renowned filter, RReliefF [6], is conducted by using Extreme Learning Machine (ELM) [28].

### 6.1. Data

Six datasets from the UCI machine learning repository [48] and the Regression website [47] were used in the experiments:

1. Abalone: the goal is to predict the age of marine snails (the abalones) from physical measurements. This dataset is available from both the UCI machine learning repository and the Regression website. It originally contains 4177 instances and 8 features. Among the instances, two outliers were detected and removed (instances 1418 and 2052) and only the physical measurements were considered. Consequently, the resulting dataset consists of 4175 instances characterized by 7 features.
2. Ailerons: this dataset can be found on the Regression website. It contains 13750 instances and 40 continuous features describing the status of a F16 aircraft. The goal is to predict the control actions on the ailerons. Some issues affect the last 16 features which can be treated as either nominal or continuous. Since the decision may have an influence on the results of RReliefF, these input variables were not considered and the dataset used in the experiments consists of 13750 instances and 24 features.
3. Boston Housing: The objective is to predict the housing prices in areas of Boston. The dataset contains 506 instances and 13 features. It is available from both the UCI machine learning repository and the Regression website.
4. CompAct: The goal is to predict the portion of time (%) during which a set of CPUs run in user mode. The prediction is performed using a collection of computer system activity measures. The data set con-

tains 8192 instances and 21 features. It is available on the Regression website.

5. Parkinson’s Telemonitoring: One of the objectives is to predict the motor score of the Unified Parkinson’s Disease Rating Scale (UPDRS) from 16 biomedical voice measures. Telemonitoring devices were used to automatically capture speech signals in 42 patients’ homes and 5875 voice recordings were made available. One of the voice measures was removed, since it provides the same value for all the recordings and the final data set used in the experiments consists of 5875 instances and 15 features. It can be downloaded from the UCI machine learning repository.
6. CT slice: the task is to predict the relative location of Computer Tomography (CT) slices on the axial axis of the human body. The prediction is carried out using features extracted from CT images. The dataset originally contains 53500 instances described by 385 features. But, 63 instances are replicated several times and 5 features provide a constant value. Consequently, the data preprocessing resulted in a slightly modified dataset consisting of 53437 instances and 380 features. The full data set is available on the UCI machine learning repository.

## 6.2. Feature Selection with MBFR

The MBFR algorithm was applied to the real world datasets described in the previous subsection. For each of them, the set of values of the parameter  $\ell^{-1}$  was chosen as follows:

1. The plot relating  $\log_e(I_{m,\ell})$  to  $\log_e(\ell^{-1})$  was computed for the full dataset (including all the instances, all the features and the target variable) with  $\ell^{-1} \in \{1, 2, \dots, 130\}$ .
2. The upper and lower bounds of the set were given by the extent of the linear part of the plot. For most of the datasets, the upper bound turned out to be simply the maximum value of  $\ell^{-1}$  ensuring the presence of two points in, at least, one quadrat.
3. If the upper bound was lower than 30, every integer value within the bounds was retained. But, if it was equal to or higher than 30, only the integer values following a geometric progression with ratio 2 were used and the bounds were modified accordingly. This allows the MBFR algorithm to run faster by reducing the value of  $R$  (see Section 3).

Datasets	Parameter $\ell^{-1} \in$
Abalone	$\{4, 8, 16, 32, 64\}$
Ailerons	$\{5, 6, \dots, 24, 25\}$
Housing	$\{2, 3, \dots, 18, 19\}$
CompAct	$\{1, 2, \dots, 9, 10\}$
Parkinson	$\{1, 2, 4, 8, 16, 32, 64\}$
CT slice	$\{1, 2, 4, 8, 16, 32, 64, 128\}$

Table 5: Sets of values of the parameter  $\ell^{-1}$  used throughout the SFS procedure of MBFR

Datasets	# Sel. Feat.	$\widehat{Diss}$	$DR$
Abalone	3(7)	0.36	0.46
Ailerons	7(24)	0.22	0.71
Housing	8(13)	0.14	0.84
CompAct	6(21)	0.03	0.94
Parkinson	8(15)	0.58	0.31
CT slice	22(380)	$\approx 0.00$	$\approx 1.00$

Table 6: Results of the application of MBFR to the real world datasets. The values in brackets in the 2<sup>nd</sup> column are the total numbers of features in the datasets and “Sel. Feat.” stands for “Selected Features”. Besides,  $\widehat{Diss}$  and  $DR$  were computed by considering only the selected features.

The resulting sets of values of the parameter  $\ell^{-1}$  are given in Table 5. Although it might seem better to change these values for each feature combination, it turned out not to be necessary. Once a set had been built by following the described procedure, it was applicable throughout the feature selection process. The CompAct dataset was the only exception: the log-log plots of several single features were characterized by two distinct linear parts. The steepest one was retained, since it led to a higher value of  $DR$ . Finally, Figure 9 displays the log-log plots of the full datasets (including the target variables) constructed with the chosen parameter sets. The values of  $S_2$  and  $M_2$  derived from the red regression lines are also indicated (see Equation 4). They suggest that the dimension of the spaces in which the datasets truly reside could be smaller than that of their embedding spaces. Consequently, the datasets could contain redundant information that MBFR might uncover.

The MBFR algorithm was applied to each dataset with the parameter values of Table 5. The results are displayed in Figure 10. In each panel, a relatively clear cut-off point allows the user to identify the features to

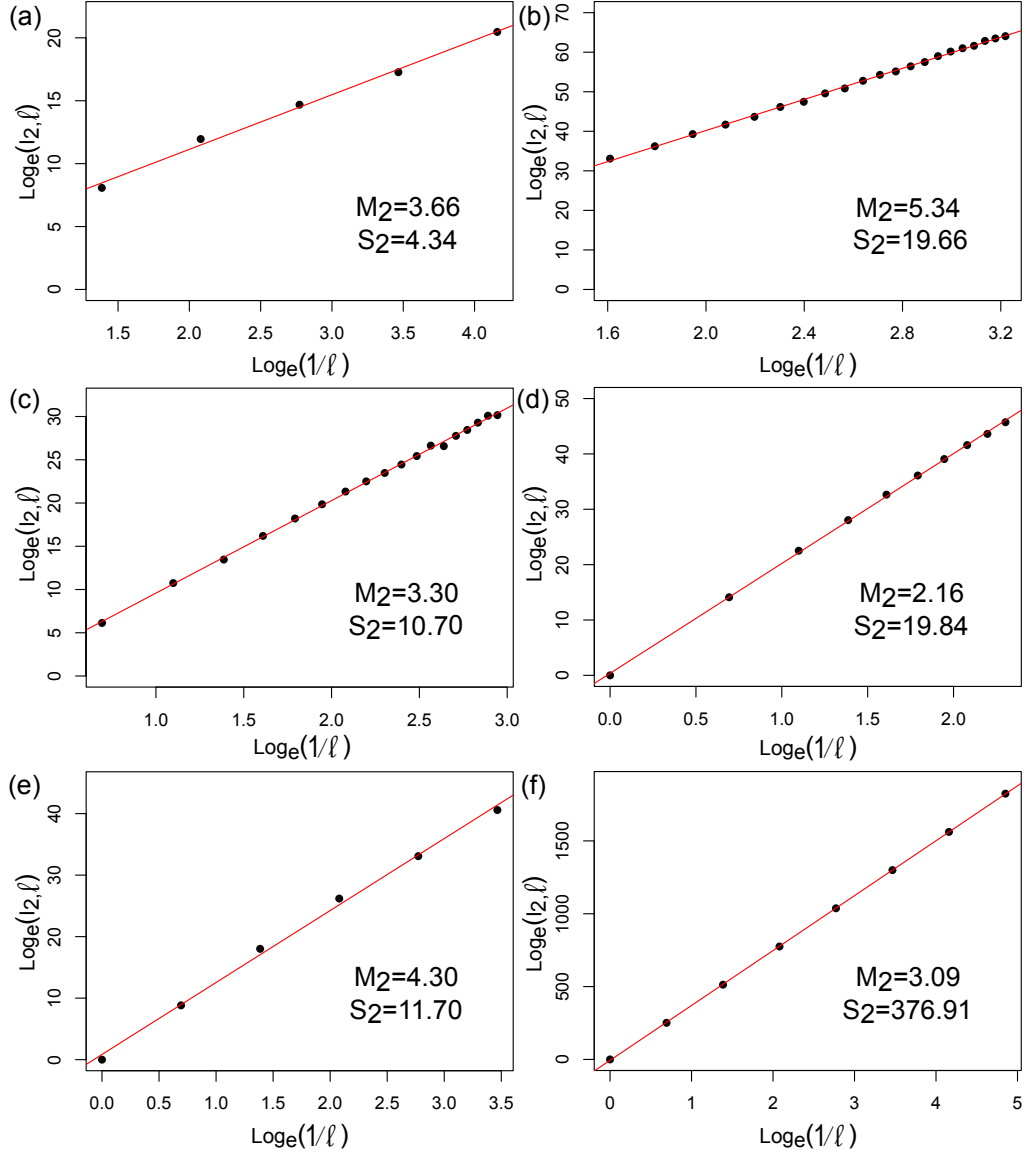


Figure 9: Intrinsic Dimension Estimation: (a) Abalone (b) Ailerons (c) Boston Housing (d) CompAct, (e) Parkinson's Telemonitoring, (f) CT slice.








Figure 11: Results of each term of Equation 9 for the Ailerons dataset. “Target” refers to the output (or target) variable of a dataset.

be selected. Moreover, Table 6 summarizes the results by providing the dissimilarity estimates (i.e.  $\widehat{Diss}$ ) and the values of  $DR$  that were computed by considering only the selected features. This overview shows that MBFR leads to a significant reduction in the number of features and that the chosen datasets cover a wide range of situations. For instance, in the Parkinson dataset, 8 features out of 15 are detected as relevant and they account for about a third of the information contained in the target variable. In contrast, in the CT slice dataset, 359 features are considered redundant or irrelevant and the relevant ones fully explain the relative location of the CT slices (i.e the target variable).

As mentioned for the simulated data, another interesting aspect of the MBFR algorithm is its ability to distinguish between redundant and irrelevant information. For instance, regarding the Ailerons dataset, Figure 11 shows that the features appearing between  $F8$  and  $F23$  in the final ranking do not contribute to the relatively high ID estimate of the data and, consequently, they can be considered redundant. In contrast,  $F5$  contains irrelevant information, but it is not fully irrelevant, since the corresponding increase in the data ID is approximately equal to 0.5 which is less than the

feature ID ( $M_2(F5) = 0.92$ ). Finally,  $F1$  and  $F2$  require further analysis. They could be either redundant with  $F5$  (and, consequently, partly irrelevant to the target variable) or simply redundant with the relevant features. To find out, the MBFR algorithm could be rerun without  $F5$  to check if the data ID is also modified by  $F1$  and  $F2$ .

In conclusion, this subsection highlights the effectiveness of the MBFR algorithm for feature selection in real world applications. In the next subsection, the results will be validated by means of a comparison with a benchmark algorithm called RReliefF.

### 6.3. Comparison with RReliefF Using Extreme Learning Machine

In this subsection, the performance of MBFR is compared with that of RReliefF [6, 49, 50] using a comprehensive evaluation procedure based on Extreme learning Machine (ELM) [28, 51]. A brief introduction to ELM and the Relief family of algorithms is also provided for a good understanding of the results.

The Relief family comprises three main algorithms of feature selection: Relief [52], ReliefF [53–55] and RReliefF [6, 49, 50]. They all consist in attributing a score to each feature according to how well its values distinguish between instances that are close to each other in the feature space. Relief achieves this goal for two-class classification problems. It randomly selects an instance and searches for its two nearest neighbours: one from the same class and one from the other class. After that, it updates the score of the features depending on their values for the randomly selected instance and the two neighbours. The operation is repeated  $m_{RF}$  times and a final score higher than 0 indicates that a feature is relevant. Following a similar procedure, the ReliefF algorithm is an extension which can deal with multiclass problems and with incomplete and noisy data.

RReliefF (i.e. Regressional ReliefF) is an adaptation of Relief and ReliefF to regression. In regression, the exact knowledge of whether two instances belong the same class or not cannot be used. RReliefF replaces it with a probability that the predicted values of two instances are different. Moreover, it computes the final score of each feature by taking into account  $k_{RF}$  neighbours. Among them, the closest ones should have a greater influence and a kernel of parameter  $\sigma_{RF}$  can be used to assign a weight to each

Datasets	RReliefF 0.00	RReliefF 0.01	MBFR	All Features
Abalone	7	1	3	7
Ailerons	24	8	7	24
Housing	10	8	8	13
CompAct	21	20	6	21
Parkinson	2	1	8	15

Table 7: Comparison between RReliefF and MBFR based on the number selected features. For comparison purposes, the last column “All Features” gives the total number of features of each dataset.

of them.

The algorithms of the Relief family have often been used as benchmarks [8, 40, 56]. In this research, RReliefF was applied to the real world datasets of Subsection 6.1 with the exception of the CT slice data that contain too many instances. The R package “CORElearn” [57] was used with the evaluation heuristic “RReliefFexpRank” and its default parameters:  $m_{RF} = N$  ( $N$  is the number of instances in the datasets),  $k_{RF} = 70$  and  $\sigma_{RF} = 20$ . These parameters were tested and turned out to be suitable for each dataset. Besides, RReliefF was run with two relevance thresholds: 0.00 and 0.01 as suggested in [8]. Any feature with a final score less than the specified threshold was considered irrelevant.

Table 7 compares the ability of RReliefF and MBFR to reduce the dimensionality of the datasets. In all but one case, MBFR performs more feature selection than RReliefF with the relevance threshold of 0.00. The differences are less pronounced with the relevance threshold of 0.01 for which the two algorithms achieve comparable results on three datasets. Nevertheless, they still provide distinct outcomes for the Compact and the Parkinson data and they never select the exact same features. Consequently, the comparison requires a way of assessing the amount of information contained in the selected features. This is the reason why ELM is used in this paper.

ELM is a single layer feed-forward neural network which can achieve the same accuracy as the well-known Multilayer Perceptrons (MLP) [58–61], while being much faster. Its main specificity lies in the weights connecting the inputs to the hidden layer. They are randomly generated and never updated and, then, the weights between the hidden layer and the outputs are learnt in a single step. Consequently, the number  $\tilde{N}$  of hidden nodes is

the only hyperparameter of the algorithm and its implementation is rather straightforward. However, an activation function is also required and a sigmoid function was used in this paper.

ELM has been applied successfully in a wide range of case studies [10, 62–64] and its high efficiency makes it well-suited to carry out comparisons. It is also worth mentioning that RReliefF and ELM have already been combined to effectively improve learning performance in temperature forecasting [65]. In this paper, ELM was used to evaluate the explanatory power of subsets of features according to a procedure which was partly presented in [66, 67]. It can be subdivided into 5 steps:

1. 20% of the  $N$  instances are randomly assigned to a testing set and the remaining 80% are passed on to Step 2. The same split is used for all the subsets of features being compared.
2. The data are projected into the  $[0, 1]$  interval and the coefficients of the projection are recorded.
3. The value of the hyperparameter  $\tilde{N}$  is selected by performing 10-fold cross-validation. For each value of  $\tilde{N} \in \{1, \dots, 350\}$ , 1 fold is iteratively allocated to a set of validation and the remaining 9 folds are assigned to a training set. For each of the 10 iterations, an ELM model is fit on the training observations and a Mean Square Error (MSE) is computed using the validation set. Then, the 10 MSEs are averaged to provide an estimate of the true error and their standard deviation is recorded. Finally, the value of  $\tilde{N}$  resulting in the lowest error estimate is retained for Step 4, unless the corresponding variability is too high. In that case,  $\tilde{N}$  is manually tuned to find a good compromise between the mean error and the variance.
4. Using all the instances involved in the cross-validation procedure (i.e. 80% of the original data), a new model is trained with the value of  $\tilde{N}$  from Step 3. Then, a prediction is made for the instances of the testing set (i.e. 20% of the original data) after they have been projected to the  $[0, 1]$  interval using the coefficient of Step 2. The operation (training and prediction) is repeated 100 times to account for the variability of the weights connecting the inputs to the hidden nodes. The final prediction for each instance is computed by averaging the 100 values and by rescaling the results to the original output range (using coefficients of Step 2). Finally, the Relative Mean Squared Error (ReMSE) of the testing set [49] is calculated. It is defined as

Datasets	RReliefF 0.00	RReliefF 0.01	MBFR	All Features
Abalone	-	0.57(0.02)	<b>0.46(0.03)</b>	0.43(0.02)
Ailerons	-	0.24(0.01)	<b>0.15(0.01)</b>	0.15(0.01)
Housing	0.17(0.06)	0.20(0.06)	<b>0.16(0.06)</b>	0.14(0.05)
CompAct	-	<b>0.02(0.00)</b>	<b>0.02(0.00)</b>	0.02(0.00)
Parkinson	0.90(0.02)	0.98(0.01)	<b>0.82(0.02)</b>	0.81(0.02)

Table 8: Comparison between RReliefF and MBFR based on the Relative Mean Squared Error (ReMSE). The table provides the mean values and the standard deviations (in brackets) of  $RE_t$  over the 20 iterations of the evaluation procedure. These values were given only if the feature selection process led to a reduction in the dimensionality of the data. Moreover, the best results (among the three feature selection approaches) are indicated in bold script.

follows:

$$RE_t = \frac{\sum_{i=1}^{N_t} (y_i - \hat{y}_i)^2}{\sum_{i=1}^{N_t} (y_i - \bar{y}_t)^2} \quad (13)$$

where  $N_t$  is the number of instances in the testing set,  $y_i$  is the measurement of the output variable for the  $i^{th}$  instance,  $\hat{y}_i$  is the corresponding prediction and  $\bar{y}_t$  is the mean of the output variable computed on the testing set. Lower values of the ReMSE are better and a value higher than 1 indicates that the tested model performs worse than the mean.

5. Steps 1 to 4 are repeated 20 times to account for the randomness in the data splits of step 1. After that, the mean and standard deviation of the 20 values of  $RE_t$  are calculated and are used to assess the explanatory power of the selected features.

For each dataset, four sets of features were passed on to the evaluation procedure: the set selected by MBFR, the two sets selected by RReliefF (one for each relevance threshold) and the full set containing all the features. Table 8 presents the results by providing the mean values of  $RE_t$  over the 20 iterations along with the corresponding standard deviations. The sets of features selected by MBFR provide comparable or better accuracy than those resulting from RReliefF. This is true even when MBFR performs more feature selection. Besides, although none of the tested algorithms improve the performance of ELM, MBFR is able to maintain or stay close to the accuracy achieved with the full datasets.

Another interesting point is the relationship between  $RE_t$  and the coefficient of dimensional relevance  $DR$ . The correlation between the two

measures was computed with the mean values of  $RE_t$  resulting from MBFR and the values of  $DR$  given in Table 6. It turned out that Pearson’s coefficient is equal to  $-0.96$ , which tends to confirm that  $DR$  is a promising measure of feature relevance.

## 7. Conclusion

This paper presents a new algorithm for supervised feature selection, namely the Morisita-Based Filter for Regression problems (MBFR). As its name suggests, it is designed for regression problems and it relies on the recently introduced Morisita estimator of Intrinsic Dimension (ID). Comprehensive numerical experiments were carried out using two simulated datasets: the well-known Friedman dataset and the butterfly dataset which was specifically designed for the needs of this research. Different sample sizes, noise levels and non-linear dependencies were tested and the variability of the results was examined by means of Monte Carlo simulations. MBFR was shown to be an effective tool for reducing the dimensionality of large datasets of varying complexity. Besides, the ability of the algorithm to distinguish between redundant and irrelevant information was presented and successfully tested.

MBFR was applied to real world datasets from publicly accessible repositories. An innovative methodology was implemented to conduct a comparison with a benchmark algorithm called RReliefF. MBFR resulted in better or comparable performance according to the accuracy achieved by Extreme Learning Machine (ELM) models. This was true even when MBFR retained fewer features than RReliefF.

A new coefficient of relevance was introduced, namely the coefficient of dimensional relevance  $DR$ . It was estimated by using MBFR and its reliability was evaluated by means of ELM models.  $DR$  is exclusively based on the ID concept, it is easily interpretable and it can be applied to high-dimensional datasets.

Finally, this paper shows that ID-based methods have a promising potential to improve the performance of existing machine learning algorithms. In addition to the presented work, they can also contribute to the development of new powerful tools to conduct fundamental tasks, such as classification, clustering and pattern detection.

## 8. Acknowledgements

The authors are grateful to Mohamed Laib and Zhivko Taushanov for many fruitful discussions about machine learning and statistics.

## References

- [1] R. Bellman, *Adaptive Control Processes: A Guided Tour*, Princeton University Press, Princeton (US-NJ), 1961.
- [2] I. Guyon, A. Elisseeff, An introduction to variable and feature selection, *Journal of Machine Learning Research* 3 (2003) 1157–1182.
- [3] I. Guyon, S. Gunn, M. Nikravesh, L. A. Zadeh, *Feature Extraction: Foundations and Applications*, Springer, Berlin, 2006.
- [4] H. Liu, H. Motoda, *Computational Methods of Feature Selection*, Chapman & Hall/CRC, 2007.
- [5] I. A. Gheyas, L. S. Smith, Feature subset selection in large dimensionality domains, *Pattern Recognition* 43 (1) (2010) 5–13.
- [6] M. Robnik-Šikonja, I. Kononenko, Theoretical and Empirical Analysis of ReliefF and RReliefF, *Machine Learning* 53 (1) (2003) 23–69.
- [7] H. C. Peng, F. Long, C. Ding, Feature selection based on mutual information criteria of max-dependency, max-relevance, and min-redundancy, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 27 (8) (2005) 1226 – 1238.
- [8] M. A. Hall, Correlation-based feature selection for discrete and numeric class machine learning, in: *Proceedings of the 17th International Conference on Machine Learning (ICML)*, Stanford (USA), 2000.
- [9] R. Kohavi, G. H. John, Wrappers for feature subset selection, *Artificial Intelligence* 97 (1-2) (1997) 273–324.
- [10] M. Leuenberger, M. Kanevski, Feature selection in environmental data mining combining simulated annealing and extreme learning machine, in: *Proceedings of the 22nd European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (ESANN)*, d-side pub., 2014, pp. 601–606.
- [11] R. Tibshirani, Regression shrinkage and selection via the Lasso, *Journal of the Royal Statistical Society B* 58 (1) (1996) 267–288.
- [12] L. Breiman, Random forests, *Machine Learning* 45 (1) (2001) 5–32.
- [13] S. F. Cotter, K. Kreutz-Delgado, B. D. Rao, Backward sequential elimination for sparse vector subset selection, *Signal Processing* 81 (9) (2001) 1849–1864.
- [14] S. Colak, C. Isik, Feature subset selection for blood pressure classification using orthogonal forward selection, in: *Proceedings of the 29th IEEE Annual Bioengineering Conference*, 2003, pp. 122–123.
- [15] X. Wang, J. Yang, X. Teng, W. Xia, J. Richard, Feature selection based on rough sets and particle swarm optimization, *Pattern Recognition Letters* 28 (4) (2007) 459–471.
- [16] S. Kirkpatrick, C. D. Gelatt, M. P. Vecchi, Optimization by simulated annealing, *Science* 220 (4598) (1983) 671–680.



- [17] R. Meiri, J. Zahavi, Using simulated annealing to optimize the feature selection problem in marketing applications, *European Journal of Operational Research* 171 (3) (2006) 842–858.
- [18] J. Kennedy, R. Eberhart, Particle swarm optimization, in: *Proceedings of the IEEE International Conference on Neural Networks*, Vol. 4, 1995, p. 1942–1948.
- [19] A. E. Hoerl, R. W. Kennard, Ridge regression: Biased estimation for nonorthogonal problems, *Technometrics* 12 (1) (1970) 55–67.
- [20] D. F. Specht, A general regression neural network, *IEEE Transactions on Neural Networks* 2 (6) (1991) 568–576.
- [21] S. Robert, L. Foresti, M. Kanveski, Spatial prediction of monthly wind speeds in complex terrain with adaptive general regression neural networks, *International Journal Climatology* 33 (7) (2013) 1793–1804.
- [22] J. Golay, M. Leuenberger, M. Kanevski, Morisita-based feature selection for regression problems, in: *Proceedings of the 23rd European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (ESANN)*, d-side pub., 2015, pp. 279–284.
- [23] F. Camastra, Data dimensionality estimation methods: a survey, *Pattern Recognition* 36 (12) (2003) 2945 – 2954.
- [24] J. A. Lee, M. Verleysen, *Nonlinear Dimensionality Reduction*, Springer, New-York, 2007.
- [25] F. Camastra, A. Staiano, Intrinsic dimension estimation: Advances and open problems, *Information Sciences* 328 (2016) 26 – 41.
- [26] J. Golay, M. Kanevski, A new estimator of intrinsic dimension based on the multipoint Morisita index, *Pattern Recognition* 48 (12) (2015) 4070–4081.
- [27] C. Traina Jr., A. J. M. Traina, L. Wu, C. Faloutsos, Fast feature selection using fractal dimension, in: *Proceedings of the XV Brazilian Symposium on Databases (SBBD)*, 2000, p. 158–171.
- [28] G.-B. Huang, Q.-Y. Zhu, C.-K. Siew, Extreme learning machine: Theory and applications, *Neurocomputing* 70 (1-3) (2006) 489–501.
- [29] B. Eriksson, M. Crovella, Estimation of intrinsic dimension via clustering, Tech. rep., Boston University, Department of Computer Science (2011).
- [30] J. Golay, M. Kanevski, C. D. Vega Orozco, M. Leuenberger, The multipoint Morisita index for the analysis of spatial patterns, *Physica A* 406 (2014) 191–202.
- [31] P. Grassberger, I. Procaccia, Measuring the Strangeness of Strange Attractors, *Physica D* 9 (1-2) (1983) 189–208.
- [32] S. Borgani, G. Murante, A. Provenzale, R. Valdarnini, Multifractal Analysis of the Galaxy Distribution: Reliability of Results from Finite Data Sets, *Physical Review E* 47 (6) (1993) 3879–3888.
- [33] S. Lovejoy, D. Schertzer, A. Tsonis, Functional Box-counting and Multiple Elliptical Dimensions in Rain, *Science* 235 (4792) (1987) 1036–1038.
- [34] Q. Huang, J. R. Lorch, R. C. Dubes, Can the fractal dimension of images be measured?, *Pattern Recognition* 27 (3) (1994) 339–349.
- [35] B. B. Mandelbrot, *The Fractal Geometry of Nature*, W.H. Freeman, San Francisco, 1983.
- [36] E. Ott, *Chaos in dynamical systems*, Cambridge University Press, Cambridge (UK), 1993.

- [37] H. G. E. Hentschel, I. Procaccia, The Infinite Number of Generalized Dimensions of Fractals and Strange Attractors, *Physica D* 8 (3) (1983) 435–444.
- [38] D. Mo, S. H. Huang, Fractal-based intrinsic dimension estimation and its application in dimensionality reduction, *IEEE Transactions on Knowledge and Data Engineering* 24 (1) (2012) 59–71.
- [39] C. Traina Jr., A. J. M. Traina, C. Faloutsos, Fast feature selection using fractal dimension - Ten years later, *Journal of Information and Data Management* 1 (1) (2010) 17–20.
- [40] E. P. M. De Sousa, C. Traina Jr., A. J. M. Traina, L. Wu, C. Faloutsos, A fast and effective method to find correlations among attributes in databases, *Data Mining and Knowledge Discovery* 14 (3) (2007) 367–407.
- [41] H. D. Lee, M. C. Monard, F. C. Wu, A fractal dimension based filter algorithm to select features for supervised learning, in: J. S. Sichman, H. Coelho, S. O. Rezende (Eds.), *Advances in Artificial Intelligence - IBERAMIA-SBIA 2006*, Springer, 2006, pp. 278–288.
- [42] D. T. Pham, M. S. Packianather, M. S. Garcia, M. Castellani, Novel feature selection method using mutual information and fractal dimension, in: *Proceedings of the 35th Annual Conference of IEEE on Industrial Electronics (IECON)*, 2009, p. 3393–3398.
- [43] R Development Core Team, *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria, ISBN 3-900051-07-0 (2008).  
URL <http://www.R-project.org>
- [44] S. H. Hurlbert, Spatial Distribution of the Montane Unicorn, *Oikos* 58 (3) (1990) 257–271.
- [45] M. Morisita, Measuring of the Dispersion of Individuals and Analysis of the Distributional Patterns, *Memoires of the Faculty of Science (Serie E)*, Kyushu University 2 (4) (1959) 215–235.
- [46] J. H. Friedman, Multivariate Adaptive Regression Splines, *The Annals of Statistics* 19 (1) (1991) 1–67.
- [47] L. Torgo, *Regression DataSets*.  
URL <http://www.dcc.fc.up.pt/~simlorgo/Regression/DataSets.html>
- [48] M. Lichman, *UCI Machine Learning Repository*.  
URL <http://archive.ics.uci.edu/ml>
- [49] M. Robnik-Šikonja, I. Kononenko, Context sensitive attribute estimation in regression, in: *Proceedings of the ICML-96 Workshop on Learning in Context-Sensitive Domains*, Bari(IT), 1996, pp. 43–52.
- [50] M. Robnik-Šikonja, I. Kononenko, An adaptation of Relief for attribute estimation in regression, in: *Proceedings of the 14th International Conference on Machine Learning (ICML)*, Nashville (USA), 1997, pp. 296–304.
- [51] G. Huang, G.-B. Huang, S. Song, K. You, Trends in extreme learning machines: A review, *Neural Networks* 61 (2015) 32–48.
- [52] K. Kira, L. A. Rendell, The feature selection problem: Traditional methods and a new algorithm, in: *Proceedings of the 10th national conference on Artificial intelligence*, San Jose (US-CA), 1992, pp. 129–134.
- [53] I. Kononenko, Estimating attributes: analysis and extensions of relief, in: F. Bergadano, L. De Raedt (Eds.), *Machine Learning: ECML-94*, Springer, 2006,

- pp. 171–182.
- [54] I. Kononenko, E. Šimec, M. Robnik-Šikonja, Overcoming the myopia of inductive learning algorithms with RELIEFF, *Applied Intelligence* 7 (1) (1997) 39–45.
  - [55] W. Duch, Filter methods, in: I. Guyon, S. Gunn, M. Nikraves, L. A. Zadeh (Eds.), *Feature Extraction: Foundations and Applications*, Springer, Berlin, 2006, pp. 89–117.
  - [56] N. Hoque, D. K. Bhattacharyya, J. K. Kalita, Mifs-nd: a mutual information-based feature selection method, *Expert Systems with Applications* 41 (14) (2014) 6371 – 6385.
  - [57] M. Robnik-Šikonja, P. Savicky, J. Adeyanju Alao, CORElearn: Classification, Regression and Feature Evaluation, R package version 0.9.45 (2015).  
URL <http://CRAN.R-project.org/package=CORElearn>
  - [58] F. Rosenblatt, The perceptron: a probabilistic model for information storage and organization in the brain, *Psychological Review* 65 (6) (1958) 386–408.
  - [59] F. Rosenblatt, *Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms*, Spartan, Washington (DC), 1962.
  - [60] P. Werbos, Beyond regression: New tools for prediction and analysis in the behavioral sciences, Ph.D. thesis, Harvard University (1974).
  - [61] D. Rumelhart, G. Hinton, R. Williams, Learning internal representations by error propagation, in: D. Rumelhart, J. McClelland (Eds.), *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, Vol. 1, The MIT Press, Cambridge (USA), 1986, p. 318–362.
  - [62] Z. H. You, Y. K. Lei, J. F. Xia, B. Wang, Prediction of protein-protein interactions from amino acid sequences with ensemble extreme learning machines and principal component analysis, *BMC Bioinformatics* 14 (8).
  - [63] A. Baradarani, Q. M. J. Wu, M. Ahmadi, An efficient illumination invariant face recognition framework via illumination enhancement and DD-DTCWT filtering, *Pattern Recognition* 46 (1) (2013) 57–72.
  - [64] W. Li, D. Wang, T. Chai, Burning state recognition of rotary kiln using elms with heterogeneous features, *Neurocomputing* 102 (2013) 144–153.
  - [65] F. Mateo, J. J. Carrasco, M. Millán-Giraldo, A. Sellami, P. Escandell-Montero, J. M. Martínez-Martínez, E. Soria-Olivas, Temperature forecast in buildings using machine learning techniques, in: *Proceedings of the 21st European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (ESANN)*, d-side pub., 2013, pp. 357–362.
  - [66] A. Jain, D. Zongker, Feature selection: Evaluation, application, and small sample performance, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 19 (2) (1997) 153–158.
  - [67] J. Reunanen, Overfitting in making comparisons between variable selection methods, *Journal of Machine Learning Research* 3 (2003) 1371–1382.